


SPACE SCIENCES LABORATORY



GPO PRICE \$ _____


CFSTI PRICE(S) \$ _____

Hard copy (HC) 1.00

Microfiche (MF) 1.50

ff 653 July 65

UNIVERSITY OF CALIFORNIA
BERKELEY CALIFORNIA



N66 34069

FACILITY FORM 602

(ACCESSION NUMBER)

22

(PAGES)

CR-77105
(NASA CR OR TMX OR AD NUMBER)

(THRU)

1

(CODE)

(CATEGORY)

05

ON LARGE MODELS OF SYSTEMS

C. West Churchman

Internal Working Paper No. 39

June 1966

This research was supported in part by the National Aeronautics and Space Administration under General Grant #NSG-243-62 under the University of California.

Space Sciences Laboratory
Social Sciences Project
University of California
Berkeley

ON LARGE MODELS OF SYSTEMS*

C. West Churchman

This paper is concerned with man's ability to understand large social systems--governments, industrial firms, universities, hospitals, and the like. Specifically, it deals with the attempt to rationalize such systems.

While the efforts of mid twentieth century to rationalize systems are not new, they are surely larger in scope and ambition than ever before. By "rationalizing a system" I mean the attempt to find out why the system exists--i.e., its objectives--and the optimal way in which it can attain its goals.

The basic process of rationalizing is very simple. One searches for a central quantitative measure of system performance, which has the characteristic that the more of this quantity the better. The more profit a firm makes the better. The more qualified students a university graduates, the better. The more food we produce the better. The more speakers and the longer their speeches the better.

One takes this desirable quantity and tries to relate it to the feasible activities of the system. Each significant activity contributes to the desirable quantity in some recognizable way. The contribution, in fact, can often be expressed in a mathematical function, that maps the amount of activity onto the amount of the desirable quantity. The more sales of a certain product, the higher the profit of a firm. The more

* Paper presented at The Second Stony Brook Conference on Advances in Computing, Stony Brook, New York, June, 1966.

courses we teach, the more graduates we have. The more fertilizer we use, the more food. The more conferences we hold, the more speeches we hear.

We call the mathematics that relates activities to the desirable quantity the "objective function." The objective function is a basic ingredient of our rationalization of systems, but it is not a sufficient ingredient. It is true that American culture loves large amounts of almost anything that Americans like: cigarettes, popcorn, candy, autos, pool halls, bars, restaurants, schools, zoos and politicians. But in each case there are limits, largely arising from the fact that our cupidity is so diversified. Too many cigarettes produce a lessening of fresh air, which we want in more quantity. Too much candy lessens slimness, which we want more of. Too much of one product reduces the number of other products we can make. Too much of any speech lessens the time that others can talk--and above all we all want to talk more.

In our rationalizations we handle these conflicting cupidities by constraint equations. We now say that we want to maximize the objective function, subject to a set of constraint equations. This way of stating our problem fits into an old fashioned mathematics most prominently associated with the name of Lagrange. But in size we have gone far beyond anything that mathematicians of a century ago ever envisioned. According to one account, we can handle a particular mathematical system with a million variables and thirty thousand constraint equations (by means of what is called the decomposition algorithm). Applied mathematics of today can be characterized by the Texan syndrome: it may not be very elegant but it is big.

We can also describe systems in many different ways: the functions can be linear or non-linear, the variables may take on discrete or continuous values, they may be deterministic or stochastic, and so on. The system

may look like a network of fantastic size, described in a dynamic fashion. When all our mathematical ingenuity fails, we can use large computers to simulate the system, and very ingenious techniques to "search for" the maxima of our functions.

All this must sound very impressive indeed, especially when one reflects where we were a scant twenty years ago. At that time there were no techniques for the solution of even modestly large mathematical systems, except by old fashioned and extremely laborious methods. Imagine trying to solve a set of linear equations with thirty unknowns by hand!

Grant that we can roughly measure the size of a mathematical system by the number of variables and equations, and that something like 20 variables and ten equations characterized 1945, and 10^6 variables and 30,000 equations characterizes 1965. If the growth is exponential, then perhaps by the time of Orwell's year 1984 we can handle systems up to 100 billion unknowns and 100 million constraint equations.

That would be quite a bit of rationalization. We should then, for example, be able to rationalize a large company like General Electric and its hundred or so divisions. Or we might even be able to rationalize a large part of the Federal government.

What can we say about this prospect? Should we look forward to a day when our government and industry will finally stop operating in a piecemeal, irrational way, and when all of the major interdependencies will be considered in a precise and rational fashion? All we would need to do, you see, is to write down what we want our federal, or local government to accomplish, and all of the reasonable constraints we wish to impose on it in terms of money, time, and other resources. The model would be written in computer code and the solution would tell each agency how it should behave.

The model could conceivably be updated daily or weekly, and there we are.

Where? Has human initiative disappeared? Is there anything left in the decision process that is human at all?

Well, of course there is still a lot of humanity engaged in the large model rationalizing. After all, humans have to write down the specifications in terms of objectives and constraints. All the models do is to rationalize what humans feel they want. And, incidentally, if the humans write down ridiculous constraints, the models can tell them this. By means of a technique often labeled as the "dual," one can judge how costly a constraint really is, and hence whether it should be relaxed or tightened.

Thus--in some age before us--we may perhaps merely say what we want and describe to a computer what we are able or willing to do, and our laws will be printed out for us. We won't understand why the results are what they are, because we won't be able to follow through the enormous computations--but we will be rationalized--won't we?

It is interesting to observe that the usual negative response to this question seems largely irrelevant. People want to say that we cannot quantify human values, nor can we put our trust in the computer. But we can and do quantify values all the time, in terms of dollars, or numbers of people, or houses, or deaths. In fact, "quantity" is simply one very ingenious and general way that man has found to describe his world, and there is really nothing basically wrong with this mode of description for those who like to use it. Besides, quantitative and qualitative descriptions are often intertranslatable. And as for a lack of faith in computer calculation, this seems to be the totally wrong place to be suspicious. Granted that computers can go wrong, or more accurately can be programmed incorrectly, they are no different in this respect from other pieces of

machinery like autos, aircraft, and buildings in which we put our trust without much thought at all. Actually, more attention seems to be spent in the computer industry on safety and reliability than in most other large equipment industries.

No, the question of whether we shall put our faith in large models is not answered by either a distrust of quantity or high speed computation. The question, indeed, is a philosophical one, that belongs to the very difficult and generally neglected area of the philosophy of inquiry. It is not a new question at all. Plato in the Timaeus, and his contemporary Democritus in writings that did not survive, both seemed to have believed that it is possible to sweep the entire world into an ever expanding model. Their belief was echoed down the ages--in Spinoza's Ethics, in nineteenth century mechanism and evolutionism, and now today in a philosophy that is often used to sell system science and operations research. Nothing, it is said, can escape the eventual embrace of rational models.

The trouble with this philosophy is that it is wrong, dangerously wrong, pigheadedly wrong, philosophically inexcusable. It is not easy to say why it is wrong, because its basic idea is so attractive to narrow minded but brilliant thinking types. But if anyone will pause for a moment and reflect, he is bound to see the serious flaw.

The flaw is an unforgivable neglect of the problem of information. The philosophy I've just called wrong assumes that information is "here," or if not here can be obtained. It's like the story of a very serious minded young psychiatrist who addressed a group of patients in a mental hospital in a paternalistic fashion. "Why, my friends," he asked, "why are we all here?" "Because," a voice boomed from the back row, "we're not all there."

It is a common habit of mind to accept reality as something fixed, out there, that we can question in various ways by means of our senses, or our senses aided by instruments. Indeed, the narrow minded scientist I referred to above likes to call this reality "Nature" (he seems to shy away from calling her "Mother Nature," but he often treats her like a woman). He speaks of "states of Nature" or "moves of Nature." He assumes that Nature is a creature he can address questions to and from whom he can receive direct answers. Thus he acts as though the facts about Nature can be "plugged into" his models, and therefore the models are realistic because they simply describe the states of Nature.

But this metaphysics is utterly naive, and is the very reason the whole philosophy of the model builder is wrong. Consider, for example, a very elementary piece of "information": the cost of performing x units of an activity. How do I ask Mother Nature this question? One suggestion is to go to the records and add up the labor and material costs in terms of past labor and material prices. Is this an answer to my question? Clearly not, because it may happen that the prices paid for labor and materials in the past were the wrong prices: the activity could easily have been mismanaged. My real question is: what is the cost when the activity is carried out properly? But how do I ask Nature this question? Which one of my senses do I use to obtain a direct answer? The only reply seems to be: you don't just use your senses; you must also reason about the activity.

Consider another example. A branch bank is a subsystem of the financial system. We can surely look forward some day to modeling the financial system of the U.S.A., and hence we shall sweep in the branch banks. How? We could describe them in terms of customer requirements, and service units of various kinds, and our model would then try to optimize the opera-

tions of the branch bank. But wait a bit. What is this subsystem trying to accomplish? Let's say that a large part of its objective is the storing and retrieval of financial information. If so, is this the kind of subsystem we need to do the job? It seems ridiculous to ask people to carry around financial information and transmit it in a haphazard fashion. Perhaps the whole subsystem should be redesigned. Perhaps we could easily eliminate checks and bank notes, which are surely very antiquated information devices.

Now note that a straightforward model builder of the financial system would never have "asked Nature" any question about the appropriateness of the subsystem--he would simply have "inputted" the data he found.

Of course, we could design such questions into the model building process. We could insist that the model builder ask whether the activity has been well managed or whether the subsystem is appropriate. But how shall he get his answers? To determine whether an activity has been well managed, we need to see its relationship to other activities. To determine whether a subsystem is appropriate, we need to understand the whole system of which it is a part.

Now we can begin to see what is wrong with the claim that large models can sweep in all reality. The models don't mean anything unless they use the correct information. But we can't determine what information is correct unless we understand how the subsystems ought to be interrelated, or unless we understand the whole system. But this is what our realistic models are supposed to tell us. In other words, we need realistic information to start with in order to build our models, but we need the model in order to get the information.

In case there is some doubt about the seriousness of questions about subsystems, let me repeat some very reasonable questions about large systems that reasonable people are asking today:

- (a) do we really need private automobiles and freeways?
- (b) do we really need a very rapid transportation system?
- (c) do we really need separate schools and universities?
- (d) do we really need libraries?

We can't go to Nature with all these "do we really's" unless we understand what the whole world is like--because "Nature" doesn't give answers to these questions in a simple, unequivocal manner.

Suppose we state the same problem in another way. If we are trying to rationalize a system, we must consider that part of the system which is capable of implementing new plans. No matter how finely we construct our model, if the decision makers won't go along with the model's conclusions, the whole exercise is meaningless except as a purely game-like mathematical exercise. Now what questions can we address to Nature the answers to which will tell us how we can be confident of implementation? The truth of the matter is that none of us know or even have a fairly good idea of the questions. Do we need to know how to "sell" new ideas? Do we need to know how to "communicate" better? Do we need to sweep politics into our models? What is an optimal implementing system? Lots of people are trying to answer these questions today--but they differ in their answers, and none of them can justify what they say in terms of "facts" gleaned from Nature's responses.

The point seems to be that we can't get "facts" about systems without making very strong assumptions about the systems.

Thus it looks as though we're involved in a vicious circle. We must get information to make our models realistic, but we must have general models to get our information.

Now there's a very direct and practical way to answer all these philosophical questions. The answer is: be practical! All this talk about the realism of the large model is fine in the abstract, but we've got to begin somewhere--so let's begin. The thing to do is start somewhere--with the feasible, and let experience modify our wrong guesses. Do something, as long as it works.

This seems to be the popular philosophy of the system science enthusiasts in the political arena. For example, a recent study was conducted to come up with a proposal for an information system for the state of California. Now the proposal certainly wasn't a full-fledged model, nor did it consider such questions as, "Do we really need a Department of Motor Vehicles," or "Do we really need educational records?" But it did say something that could be the basis of a beginning. The same practical philosophy seems to pervade the application of cost-effectiveness analysis, or PERT, or any of the other techniques. They're all merely starting points--some things we can do now.

The trouble with this pragmatic philosophy is that it's correct, absolutely correct--mainly because it doesn't say anything at all. What it says is "do what is feasible." Now if you don't know what "feasible" means, look it up and find that it means "capable of being done." So this very wise and practical philosophy says "do what is capable of being done." In some ways I prefer the opposite philosophy "do what is incapable of being done."

The problem, of course, is to find out what can be done that won't ruin us--take us down the irreversible pathways of self-destruction. If you try to pin the practical philosopher down to what he means by "feasible," he becomes very elusive, if not downright abstract. Maybe he'll point out that in mathematical programming, a "feasible" solution is one that meets all the constraints. But he doesn't mean this when he talks about the feasible; for one thing, "doing nothing at all" is often a mathematically feasible solution, but no one in his right mind would say that a do-nothing program is feasible.

No, what our practical, hard nosed realist means by "feasible" is a plan that "people will accept." He points out that many plans based on elaborate models are not feasible because no one will understand them, and people will resist them because they feel threatened.

The trouble is that once the practical philosopher moves away from his favorite tautology ("do only what you can do") he gets into deep waters he himself doesn't understand. Politics is a subsystem of every large system. Resistance to suggestions is a part of politics. If I say, "propose what people will not oppose," I'm telling you how to conduct your politics, and my advice may be very bad indeed. Sometimes politicians who try to avoid controversy are very bad politicians. The so-called practical philosophy of the feasible may really be based on a naive or even stupid political premise.

Besides, I know enough about people to know how they love to deceive themselves. How many faculty meetings I've sat through that have ended by a chairman's remarking "I think we can all agree--"! and on so many of these occasions what follows the announcement is the most unbelievable series of

recommendations imaginable! We can all agree to do the stupidest things, especially since we all agree!

No, I don't think we should base our planning on what is feasible at all, simply because this kind of down-to-earth practical philosophy is so much up in the air without any guarantee that it will fly.

What then? Should we abandon the attempt to rationalize human systems by large models? Certainly not. I don't like the "intuitive" side any better than the zealous rationalists. The intuitives like to say proudly and with textual inaccuracy, "our company just grew like Topsy," but "we've learned to fly by the seat of our pants" and "by a lot of work we've pulled ourselves up by our own bootstraps." The picture of a heavily booted and topless Topsy driven by a turbine engine in her bottom, is probably typical of the conscious confusion that goes on in the mind of the managerial intuitive. There is less and less excuse for an ignorance about modern analytical technology on the part of today's top managers in industry and government. They are irresponsible if they pretend that the use of models and computers in their planning is "beyond them," or that they don't need to know about these newer developments because they've gotten so far without them. They are equally irresponsible if they expect to see "positive" results from planning models in one or two years at a minimum expenditure of time and effort. By this late date, intuitive managers should be realizing that an understanding of how their organizations really work is at least as difficult as an understanding of how a high class rocket works. We live in an age of model building for decision making, and we can make this age the most significant of all time if we all work on the problem together.

After having thrown out the pure model builder, the pragmatist and the intuitive, what have I got left? One great asset of the human race:

disagreement. Controversy.

I believe that our fondness for being right and not being contradicted has led us into the acceptance of one of the three philosophies I've just discussed. The model building rationalizer sees the world to be a world that he can adequately and precisely describe mathematically. The practical philosopher sees the world to be the world of action--of compromise and doing. The intuitive sees the world to be one of his own making--made out of his genius. Well, each of these world makers tells us a story. He tells us what the world is really like and how we should cope with it. If our tribe were a smaller one, we could imagine that these three myth builders might each tell us the story of the future world in his own way. Each would have his heroes and his bad guys. The heroes of the rationalist are brilliant men, scientists who can carve fine images of reality and in the magic flash of colored lights, have their machines spell out what the next worlds will be like if we do certain things today. The villains of this story are irrational piece-mealers, selfish, shortsighted, dull of wit, endowed only with power.

The heroes of the practical philosophers are doers--they get people to adopt their ideas--to change a little bit here and there--in one spin-off or another--towards a better technological world. Their villain is the man who wrote on the wall of the Faculty Club at the University of California: whenever it is not necessary to change, it is necessary not to change. Their heroes are simple folk, who can talk sense to managers, who get things done, who accomplish change.

The heroes of the intuitives are the great leaders of government and industry, the fine grey haired stern fathers who grace our boards of trustees,

the judiciary, the professorial. They are wise and quick to respond when speed is essential, deliberate when deliberation is called for. Their villains are self-seeking interest groups, or dissidents of one kind or another who don't recognize their proper roles.

All these story tellers tell us fascinating tales, of horror or joy, of success or calamity. And whom shall we believe and have trust in? Who has the real insight? Which myth maker has his ear tuned to God?

Why none of them and yet all of them. We are the listeners, and if we listen well we shall hear the differences as well as the sameness. It is not necessary for us folk of the tribe to believe wholeheartedly in what model builders say, any more than we need believe in practical men or wise men. But we should listen most carefully to the story that each has to tell.

And this brings me back to the theme of the large model. A large planning model is a story--it is one idea of what reality is like and what it could be like. It is a marvelously told story in its way--not dramatic perhaps, but as a mosaic of details it is unsurpassed. One can wander endlessly in the ramifications of the fabric of the tale, touching on this or that episode and the way it will affect our lives.

The main trouble with this type of story telling is that the story tellers believe they must be consistent. Now no story teller who is worth anything at all as a concocter of tales should ever try to be consistent. Certainly the man of action regards petty consistency as an anathema, and the wise father-figures only use it as a political device when it suits their aims. There is no reason why all model builders have to tell the same story about the same system. This makes them very dull people.

I've a specific suggestion. We would all like to know whether the U.S.A. should stay in Vietnam. Some men of action tell us it's the only feasible course because how can we pull out? Some other men of action tell us it's the wrong course because we must pull out. Some wise men, stroking their beardless faces, say we must stem the tide. Other wise men say we must face the music. What do model builders say? Nothing. They are too shy. They are afraid that their "information" or forecasts may be inadequate. Nonsense. Let us build two model building story tellers. One will tell us what the world is really like and his model world will interpret all the data to show that our current policy in Vietnam is correct. The other will tell us what the world is really like and his model world will interpret all the data to show that our current policy is incorrect. The two worlds these story tellers will build are different, of course. It is up to us the listeners to see which one--if either--we are willing to accept.

I'd suggest we try this model building controversy in all kinds of contexts--educational planning, poverty wars, health, urban development, and so on. After all, debate has long been the common practice of practical men and intuitive men--why not of rationalizers as well? Actually, rationalizers do debate among each other fiercely, as any one knows who's worked in a lively team of operations researchers. One member of the team will push for one viewpoint, and find himself strongly opposed by his colleagues. The trouble is, we hide all this debate when we make our briefings. We think we should make one united, consistent proposal based on one model. The decision maker, therefore, loses out on the really crucial part of our study, the conflict of ideas that went into our model building.

But what of truth? What is the correct answer to our pressing problems of human systems? Well, there was one American philosopher, John Dewey, who showed that the quest for certainty is bound to fail, because certainty is an unattainable goal. There was another American philosopher, Edgar Singer, who put the matter more deeply and accurately. When we reach a conclusion after having exposed our ideas to the most severe test we can imagine, then we have done the best that inquiry can possibly accomplish.